Massively parallel read mapping on GPUs with PEANUT

Johannes Köster* Sven Rahmann

Genome Informatics, Institute of Human Genetics
University Hospital Essen, University of Duisburg-Essen

March 10, 2014

arXiv:1403.1706v1 [cs.DS] 7 Mar 2014

We present PEANUT (ParallEl AligNment UTility), a highly parallel GPU-based read mapper with several distinguishing features, including a novel q-gram index (called the q-group index) with small memory footprint built on-the-fly over the reads and the possibility to output both the best hits or all hits of a read. Designing the algorithm particularly for the GPU architecture, we were able to reach maximum core occupancy for several key steps. Our benchmarks show that PEANUT outperforms other state-of-the-art mappers in terms of speed and sensitivity. The software is available at `http://peanut.readthedocs.org`.

1 Introduction

A key step in many next generation sequencing (NGS) data analysis projects, e.g. for variant calling or gene expression analysis, is mapping the obtained DNA reads to a known reference sequence. We distinguish between the *read mapping problem*, which is to find the possible origins of each read within the reference, and the *read alignment problem*, which is to provide basepair-level alignments between each read and each originating region. The read mapping algorithm has to be error tolerant because of mutations and sequencing errors. The optimal solution to this problem is to calculate semiglobal optimal alignments of each read against the reference. Since this would incur an infeasible running time (proportional to the product of reference length and total read lengths, e.g. 6 Gbp $\cdot\, 10^7 \cdot 100$ bp $= 6 \cdot 10^{18}$ bp^2 for the human genome and its reverse complement against 10 million Illumina reads), various filtering methods and approximations have been developed. These can be roughly classified into methods based on backward search using the Burrows-Wheeler transform (BWT), e.g. BWA [11] or Bowtie 2 [9], and methods based on q-gram indexes, e.g. RazerS 3 [23].

An extension of the read mapping problem is the *split read mapping* problem to also find partial matches of a read, especially of the prefix and suffix of a read. Split read mapping is useful to detect differential transcript splicing (with the help of reads that span an intron) or structural variations [16]. One strategy for split read mapping is to partition the read (as done by Tophat [22] and others) and map each part with small error tolerance. If a read mapper is capable of high error tolerance, calculating local alignments can also yield split read mappings (as done by BWA-MEM [10] and others).

*To whom correspondence should be addressed.

Read mappers can be categorized [23] into best-mappers that try to find the (or any) best origin of a read (e.g. BWA-MEM) and all-mappers that provide a comprehensive enumeration of all possible locations (e.g. RazerS 3 or the "all" mode of Bowtie 2) up to a given error threshold. While all-mappers can be much slower (depending on the number of hits), their strategy is advantageous when a confidence value is required that the reported origin is the true origin of the read. Further, all-mappers are useful when mapping to homologous sequences like alternative transcripts [20] or meta-genomes, where more than one hit is expected. An intermediate strategy is to report all hits of the best stratum, i.e., all hits with the same lowest error level (instead of only the first or a random such hit).

Recently, exploiting the parallelization capabilities of graphics processing units (GPUs) for read mapping has become popular and GPU-based BWT-read-mappers appeared, e.g. SOAP3 [13] and SOAP3-dp [14]. Using a q-gram index on a GPU is not a common choice because of its large size. Therefore, to the best of our knowledge, q-gram index based mappers so far only use the GPU for calculating the alignments and keep the index on the CPU, e.g. NextGenMap [21] and Saruman [3].

Here we present PEANUT (ParallEl AligNment UTility), which is a GPU-based split read mapper using a q-gram index. PEANUT provides the first feasible implementation of a q-gram index on a GPU, which is achieved by introducing the *q-group index*, a novel q-gram index implementation with a smaller memory footprint, and by indexing subsets of the reads on-the-fly on the GPU. PEANUT is the first GPU-based all-mapper. With both a recent and a four years old NVIDIA™ Geforce GPU, we show that PEANUT outperforms other state of the art best-mappers in terms of speed and sensitivity. Futher, PEANUT is faster than other all-mappers, while achieving a comparable sensitivity.

This article is structured as follows. We first lay the foundations for the PEANUT algorithm in Section 2 by reviewing requirements for efficient GPU usage and introducing the q-group index data structure. In Section 3, we give an overview of the PEANUT algorithm and describe its different steps (filtration, validation, postprocessing) in detail. Section 4 shows several benchmark results on speed and sensitivity of PEANUT. A brief discussion concludes the paper.

2 Foundations

PEANUT uses the filtration plus validation approach: It first quickly detects exact matches of a given length q between each read and each reference (so-called q-gram matches) and only computes alignment scores where such matches are found. For increasing values of q, this substantially reduces the computational burden. The advantage of using a q-gram index (instead of a compressed suffix array or FM index) is its simplicity and quick lookup speed. The main disadvantage is its high memory usage; we overcome this by indexing a part of the reads on the fly instead of indexing the genome. We propose a novel implementation of the q-gram index functionality (the *q-group-index*) with a smaller memory footprint that is well suited for modern GPU architecture. We first discuss the GPU architecture and its implications for designing the PEANUT algorithm to maximize parallel GPU usage. Then, we describe the *q-group index* data structure.

2.1 Designing for efficient GPU usage

We use the terminology of NVIDIA™, while the general concepts are also applicable to competitors like AMD™. A GPU is partitioned into Symmetric Multiprocessors (SMs), each of which has its own on-chip memory, cache and processing cores. By adjusting the *thread block size* it can be controlled how threads are distributed among the SMs. Each SM executes one thread block until all threads in the block are completed. Once an SM has completed a thread block, it moves to the next if any blocks are left. An SM can execute 32 threads in parallel (restricting the thread

block size to be a multiple of 32); such a group of threads is called a *warp* or *wavefront*. At any time, each of these threads has to execute the same instruction in the code, but may do so on different data, a concept which is called *single instruction, multiple threads* (SIMT). Hence, conditionals with diverging branches should be avoided, since threads taking an if-branch have to wait for threads taking the corresponding else-branch to finish and vice versa. All SMs may access a slow common global memory (about ≤ 3 GB on most of today's GPUs) in addition to their fast on-chip cache and memory. While the size of the fast cache is extremely limited, accesses to the global memory are slow and should be minimized. However, the memory latency can be reduced by *coalescing* the access, i.e., letting threads in a warp access contiguous memory addresses, such that the same memory transaction can serve many threads. In addition, an SM can execute a different warp while waiting on a transaction to finish, thereby hiding the latency. For the latter, threads should minimize their register usage such that the number of warps that can reside on an SM is maximized. Finally, data transfers from the main system memory to the GPU global memory are comparatively slow. Hence it is advisable to minimize them as well.

A useful programming pattern that is used extensively in the algorithms presented in this work are *parallel prefix scans*, a special case of which is the computation of a cumulative sum, which at first appears to be a serial process, or filtering an index set. Parallel prefix scans are used to solve these problems in a data parallel way with a minimum amount of branching, thus nicely fitting the above considerations; see [2] for a comprehensive introduction. We use the PyOpenCL implementation of prefix scans [7].

2.2 The q-group index

A classical DNA q-gram index of a text T (usually the reference genome, but for PEANUT a subset of the reads) stores, for each string of length q over the alphabet, at which positions in T the q-gram occurs and allows retrieving these positions in constant time per position. It is commonly implemented via two arrays that we call the *address table* A and the *position table* P.

Q-grams are encoded as machine words of appropriate size with two subsequent bits encoding one genomic letter (i.e., A = 00, C = 01, G = 10, T = 11). Unknown nucleotides (usually encoded as N) are converted randomly to A, C, G or T, and larger subsequences of Ns are omitted from the index. Hence, a q-gram needs $2q$ bits in hardware and is represented (encoded) as a number $g \in \{0, \ldots, 4^q - 1\}$. The address table provides for each (encoded) q-gram g a starting index $A[g]$ that points into the position table such that $P[A[g]], P[A[g]+1], \ldots, P[A[g+1]-1]$ are the occurrence positions of g.

Deciding about the q-gram length q entails a tradeoff between specificity of the q-grams and the size of the data structure. Array A needs 4^q integers and thus grows exponentially with q, while array P needs $|T|$ integers, independently of q. Larger values of q lead to fewer hits per q-gram that need to be validated or rejected in later stages. Further, the choice of q determines the sensitivity or error-tolerance of the search via the pigeonhole principle (q-gram lemma): only if there are $e < (n+1)/q - 1$ errors, we can guarantee that a q-gram match exists.

Idea. The idea of the *q-group index* is to have the same functionality as the q-gram index (i.e., retrieve all positions where a given q-gram occurs in contant time per positions), but with a smaller memory footprint for large q. This is achieved by introducing additional layers in the data structure. In the following, we always consider a q-gram as its numeric representation $g \in \{0, \ldots, 4^q - 1\}$.

We divide all 4^q q-grams into groups of size w, where w is the GPU word size (typically $w = 32$). The q-gram with the number g is assigned to group number $\lfloor g/w \rfloor$. Thus the i-th group is the set $G_i = \{g \mid \lfloor g/w \rfloor = i\}$ of w consecutive q-grams according to their numeric order. The set of all q-groups is $\mathcal{G}_q := \{G_0, G_1, \ldots, G_{\lceil \frac{4^q}{w} \rceil - 1}\}$. We write g_{ij} for the j-th q-gram in G_i.

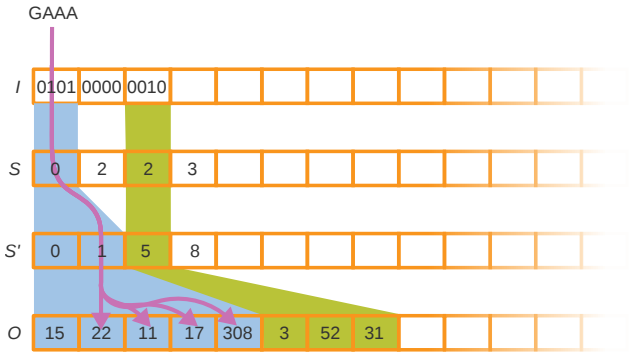

Figure 1: The q-group index consists of four arrays I, S, S', O. The purple arrows illustrate how the four layers of the index are traversed to reach the occurrences of the queried q-gram GAAA.

For a given q and text T, the q-group index is a tuple of arrays

$$\mathcal{I}_{T,q} := (I, S, S', O).$$

Array I consists of $|\mathcal{G}_q|$ words with w bits each (overall 4^q bits), and bit j of $I[i]$ indicates wheter g_{ij} occurs at all as a substring in the text, i.e.,

$$I[i]_j = \begin{cases} 1 & \text{if } g \text{ is a substring of } T, \\ 0 & \text{otherwise.} \end{cases}$$

The array O corresponds to the position table P of a regular q-gram array: it is the concatenation of all occurrence positions of each q-gram in sorted numeric q-gram order. To find where the positions of a particular q-gram g begin in O, we first determine the group index i and the q-gram number j within the group, such that $g = g_{ij}$. With the bit pattern of $I[i]$, we determine whether q_{ij} occurs in T. If not, there is nothing else to do. If yes, i.e., $I[i]_j = 1$, we determine the j' such that bit j is the j'-th one-bit in $I[i]$.

The address array S contains, for each group i, an index into another address array S', such that $S'[S[i] + j']$ is the starting index in O where the positions of g_{ij} can be found; see Figure 1 for an illustration. All occurrence positions are now listed as

$$O[S'[S[i] + j']], \ldots, O[S'[S[i] + j' + 1] - 1].$$

Similarly to a plain q-gram index, access is in constant time per position: For a given q-gram g, to determine (i, j) such that $g = g_{ij}$, we simply compute $i = \lfloor g/w \rfloor$ and $j = g - wi = j \bmod w$. To compute j', i.e., find how many one-bits occur up to bit j in $I[j]$, we use the population count instruction with a bit mask:

$$j' = \textsc{Popcount}(I[i] \,\&\, (2^j - 1)).$$

The population count $\textsc{Popcount}(x)$ returns the number of 1-bits in x.

Construction. Algorithm 1 shows how the index is built. The outline of of the algorithm is as follows. First, I is created from the q-grams of the text (line 2). Then, S is calculated as the cumulative sum over the population counts of I (line 6). Next, the number of occurences for each q-gram is calculated (line 10) and S' is created as the cumulative sum over these counts (line 14). Finally, the q-gram positions are written into the appropriate intervals of O (line 16).

Algorithm 1 Building the q-group index. For a q-gram g, the function GROUP-AND-BIT(g) computes (i, j) with group $i := \lfloor g/w \rfloor$ and bit $j := g \bmod w$. The function GROUPRANK(I, i, j) computes $j' := \text{popcount}(I[i] \,\&\, (2^j - 1))$, as explained in the text.

Input: a text T, machine word size w
Output: the q-group index (I, S, S', O)

```
1: initialize I with ⌈4^q/w⌉ zeros
2: for p ← 0, ..., |T| − 1 in parallel
3:     (i, j) ← GROUP-AND-BIT(q-gram at position p in T)
4:     I[i]_j ← 1
5: allocate S with space for |I| integers
6: for i ← 0, ..., |I| − 1 in parallel
7:     S[i] ← POPCOUNT(I[i])
8: S ← cumulative sum of S
9: initialize S' of length S[|I| − 1] with zeros
10: for p ← 0, ..., |T| − 1 in parallel
11:     (i, j) ← GROUP-AND-BIT(q-gram at position p in T)
12:     j' ← GROUPRANK(I, i, j)
13:     increment S'[S[i] + j'] by 1
14: S' ← cumulative sum of S'
15: Allocate O of length |T|
16: for p ← 0, ..., |T| − 1 in parallel
17:     (i, j) ← GROUP-AND-BIT(q-gram at position p in T)
18:     j' ← GROUPRANK(I, i, j)
19:     k ← next free entry in O after S'[S[i] + j']
20:     O[k] ← p
```

Each step is implemented on the GPU with parallel OpenCL[1] kernels. The cumulative sums are implemented with parallelized prefix scan operations (see Section 2.1). Importantly, the algorithm needs hardly any branching (hence maximizing concurrency) and makes use of coalescence along the reads in order to minimize memory latency. All major data structures are kept in GPU memory.

Size. To determine the size of the q-group index, we note that both I and S consist of $\lceil 4^q/w \rceil$ words, S' contains an index for each occurring q-gram and hence of $\min\{4^q, |T|\}$ words, and O is a permutation of text positions consisting of $|T|$ words. Thus the q-group index needs up to $2/w \cdot 4^q + \min\{4^q, |T|\} + |T|$ words, compared to a conventional q-gram index with $4^q + |T|$ words.

If $4^q \lessapprox |T|$, the conventional q-gram index has a small advantage because each q-gram occurs (even multiple times). In fact, assuming $w = 32$, if $4^q = \epsilon|T|$ for some $\epsilon < 1$, the size ratio between q-group index and q-gram index is $(\epsilon/16 + \epsilon + 1)|T|/((\epsilon + 1)|T|) = 1 + \frac{\epsilon}{16(1+\epsilon)}$. For $\epsilon = 1$ (or $|T| = 4^q$), this means a small size disadvantage of 3% for the q-group index.

If on the other hand, q becomes larger for fixed text size (such that q-grams become sparse), the q-group index saves memory, up to a factor of 16. In this case, write $4^q = K|T|$ for some $K > 1$. The size ratio is $(K/16 + 1 + 1)|T|/((1 + K)|T|) = (2 + K/16)/(1 + K)$, which tends to $1/16$ for large K. The break-even point is reached for $K = 16/15$.

In practice, we use $q = 16$ because a bit-encoded DNA 16-mer just fills a 32-bit word and $q = 16$ offers reasonable error tolerance and high specificity. In this regime with current GPU memory size, we use $|T| = 10^8$ (processing 100 million nucleotides at a time), so the size ratio is $K = 42.95$,

[1] https://www.khronos.org/opencl

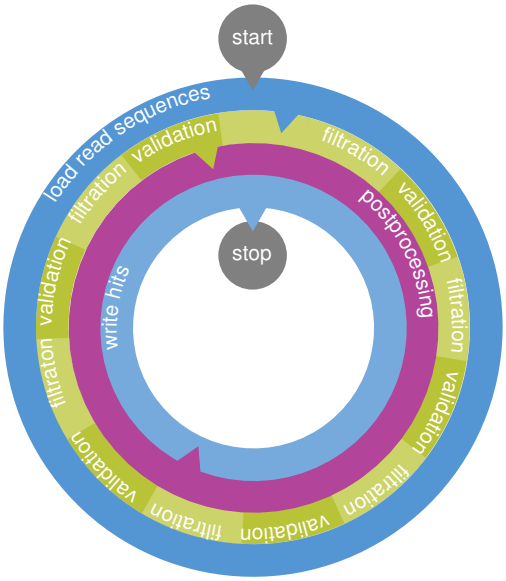
start
load read sequences
validation
filtration
validation
filtration
validation
filtration

Algorithm 2 Filtration of reference positions.

Input: reference sequence, set P of considered reference positions, maximum read length m, q-group index (I, S, S', O)

Output: array H of hits as pairs (d, r) of diagonal d and read id r

1: Initialize array C of length $|P| + 1$ with zeros to count hits
2: **for** $p \in P$ **in parallel**
3: $\quad (k_{\text{start}}, k_{\text{end}}) \leftarrow$ IndexPair(q-gram at reference position p)
4: $\quad C[p+1] \leftarrow k_{\text{end}} - k_{\text{start}}$
5: $P \leftarrow \{p \in P \mid C[p+1] > 0\}$
6: $C \leftarrow$ cumulative sum of C
7: Initialize Array H of length $2C[|P|]$ with zeros to store hits
8: **for** $p \in P$ **in parallel**
9: $\quad (k_{\text{start}}, k_{\text{end}}) \leftarrow$ IndexPair(q-gram at reference position p)
10: $\quad$ **for** $k \leftarrow 0 \ldots, k_{\text{end}} - k_{\text{start}} - 1$ **do**
11: $\quad\quad p' \leftarrow O[k_{\text{start}} + k]$
12: $\quad\quad r \leftarrow \lfloor p'/m \rfloor$
13: $\quad\quad d \leftarrow p - (p' \bmod m)$
14: $\quad\quad H[C[p] + k] \leftarrow (d, r)$

the reference and the reads.

First, a subset of reads is loaded into memory and its q-group index is created on the GPU on the fly (see Algorithm 1). If the index was built over the reference, either repetetive large data transfers or on-line rebuilding of the index would be necessary for each set of buffered reads and each chromosome. Hence – inspired by RazerS 3 [23] – we build the q-group index over concatenated reads instead of the reference, so we only build the index once for each set of buffered reads.

We now explain how to use the q-group index to retrieve potential hits. As stated above, the filtration step will be executed for a buffered set of reads, separately on each reference chromosome. The reference positions of interest are streamed against the index, and hits are obtained by reference position.

We assume that there is a function IndexPair(g) that returns, for a q-gram g, an index pair $(k_{\text{start}}, k_{\text{end}})$ such that the occurrence positions in the indexed text are all $O[k]$ with $k_{\text{start}} \leq k < k_{\text{end}}$, where O is from the q-group index built with Algorithm 1. Given the q-group index (I, S, S', O), the function IndexPair(g) is implemented as follows. Let $(i, j) :=$ Group-And-Bit(g) and $j' :=$ GroupRank(I, i, j). Then $k_{\text{start}} = S'[S[i] + j']$ and $k_{\text{end}} = S'[S[i] + j' + 1]$.

Algorithm 2 shows how putative hits are generated by querying the q-group index of buffered reads for each q-gram of the reference. First, the number of hits per reference position are computed in parallel and stored in array C (loop in line 2). In the following, only positions with at least one hit are considered (line 5). The cumulative sum of the counts generates an interval for each position, that determines where its hits are stored in the output array of the algorithm (line 6). Finally the occurences for each reference q-gram are translated into hits that are stored in the corresponding interval of the output array (loop in line 8). We translate the position inside the index of concatenated reads into a read number (line 12) and a "hit diagonal" that denotes the putative start of the read in the reference (line 13, see [19]).

Again, each step of Algorithm 2 is implemented on the GPU with parallel OpenCL kernels. The filtering of P (line 5) uses parallelized prefix scan operations (see Section 2.1). All data structures reside in GPU memory; between the steps, at most constant amounts of data have to be transferred between host and GPU (e.g., a single integer).

The set P of reference positions to investigate and the reference sequence are retrieved from a

precomputed HDF5-based reference index[2]. First, this speeds up access to the reference. Second, we can pre-compute repetitive regions (defined as exceedingly frequent q-grams) and omit them from the set P. Such a repeat masking is common in order to avoid uninformative hits when using q-gram index based algorithms. Finally, P is sorted in numerical order of the q-grams. This increases the memory coalescence when accessing the q-group index, since subsequent threads will have a higher probability to access the same region in the index and hence the same memory bank in the global GPU memory.

3.3 Validation

The validation step takes the potential hits of the filtration step and calculates the edit distance (with unit costs, i.e., the Levenshtein distance) between a read and the reference sequence at its putative mapping position. If the edit distance is small enough, the hit is considered to be correct and will be postprocessed in the next step.

The edit distance is calculated with Myers' bit-parallel algorithm [18] that simulates the edit matrix E between reference sequence and read, which contains one column for each reference base and one row for each read base. The value E_{ij} is the minimal edit distance between the read prefix of length i and any substring of the reference that ends at position j. Interpreted as a graph with a node for each matrix entry, a path in E from the top to the bottom row denotes a semi-global alignment between read and reference. Myers' algorithm calculates the edit matrix column-wise. The state of the current column is stored in bit vectors. A transition from one column to the next is achieved via a constant amount of bit-parallel operations on the bit vectors. In iteration j, the minimal distance between the read and any substring of the reference that ends at position j can be retrieved. If the accepted error rate is limited, only a part of the edit matrix is needed to calculate the optimal edit distance. Our implementation of the algorithm follows a banded version that calculates only the relevant diagonal band of the edit matrix [23]. The implementation keeps the considered part of each column in a single machine word of size w (currently 32 bits). Thereby it provides a complexity of $\mathcal{O}(|r|)$ with $|r|$ being the read length. While the reduction to the diagonal band restricts the maximum insertion or deletion size in a single alignment, mismatches are not affected. Hence, the procedure allows to discover partial matches of the read, as needed for split read mapping. Large indels can be rescued later when calculating the actual alignment if a sufficiently large portion of the read aligns in this step.

Similar to [23] we use the algorithm to calculate the edit distance of a semi-global alignment in backward direction, thereby obtaining the best starting position of the alignment. For each hit, the fraction of matches or *percent identity* is obtained from the edit distance as $(|r|-k)/|r|$, where $|r|$ is the read length and k is the edit distance. Hits with a percent identity less than a given threshold are discarded. The default for this threshold is 80 percent which provides a decent sensitivity in our benchmarks (see Section 4). Decreasing it has moderate impact on performance, since more hits have to be postprocessed and written to disk.

3.4 Postprocessing

The postprocessing of a read removes duplicate hits (generated by clusters of matching q-grams), sorts the remaining hits by percent identity (our alignment score), pairs mates in case of paired-end alignment, estimates a mapping quality and calculates the actual alignment of each hit. In contrast to the previous steps, it is done on the CPU. This allows to postprocess the hits in parallel to filtration and validation.

Intuitively, a particular hit is more likely to be the true origin of a read the fewer hits with the same or with a better score occur. During postprocessing, hits are sorted into strata of the same

[2]Hierarchical data format version 5, `http://www.hdfgroup.org/HDF5`

percent identity [15]. In paired-end mode, the percent identities of properly paired hits are summed when determining strata. Upon invocation, PEANUT can be configured to discard hits based on their stratum, e.g. providing only the best stratum or all strata. In the following, we refer to these as *best-stratum* and *all* modes.

A useful measure for distinguishing true positive hits (i.e. hits referring to the true origin of a read) from false positives is to calculate for each hit the number of hits with a percent identity at least as high. For read r with $s_{r,p}$ being the best percent identity for an alignment starting at any arbitrary reference position $p \in P$, we denote this as *hit-rank* $R_{r,p}$ and approximate it by

$$R_{r,p} = |\{p' \in P \mid s_{r,p'} \geq s_{r,p}\}| \approx |\{p' \in H_r \mid s_{r,p'} \geq s_{r,p}\}|$$

with H_r being the validated hits of read r. The hit-rank can be obtained effectively via sorting H_r by score. The intuition is that the smaller the hit-rank is (i.e. the higher the percent identity), the more likely the hit will be a true positive. In Section 4 we show that the hit-rank provides a sharp estimate to distinguish between true and false positives. An important property of the hit-rank is that it is not affected by the applied percent identity threshold since only better hits are considered for its calculation.

PEANUT provides its results in the SAM format [12], which expects a mapping quality for each hit. The mapping quality of read r at position p is expected to be the PHRED-scaled error probability, i.e., $-10\log_{10}$ of the probability that the read does not originate from position p. To obtain a mapping quality from the hit-rank, we first assume that reads are uniformly sampled from the reference. Under this null-hypothesis, the probability to obtain by chance a score X at least as extreme as $s_{r,p}$ can then be estimated by

$$\Pr(X \geq s_{r,p}) = \frac{R_{r,p} - 1}{|P|}.$$

Finally, the PHRED-scaled mapping quality is obtained as $\min\{-10\log_{10}\Pr(X \geq s_{r,p}), 255\}$.

4 Results

We evaluate PEANUT in terms of its efficiency of GPU resource usage, its sensitivity and its speed. We also evaluate the ability of the mapping quality measure defined in Section 3.4 to separate true hits from others.

GPU resource usage. In order to maximize utilization of the GPU hardware, idle cores have to be avoided. The two most important reasons for idle cores are branching and memory latency (see Section 2.1). The latter can be hidden if the SM is able to execute a different warp while waiting on the transaction. The capability to do so can be measured with the *occupancy*, that is the fraction of active warps among the maximum number of warps on an SM. The more active warps exist on an SM, the higher is the chance that latency can be hidden by executing another warp. Figure 3 shows the occupancy patterns of the implemented OpenCL kernels, as measured with the NVIDIA™ CUDA command line profiler depending on the used thread block size (see Section 2.1). The thread block size influences the occupancy by limiting the number of potentially active warps and determining the amount of used registers and shared memory on the SM. Since the latter are limited, a bigger thread block size does not necessarily lead to a higher occupancy. As can be seen, the occupancy for all steps is high. For building of the q-group index (create_queries_index) and the filtration step (filter_reference), it even reaches 1.0 which illustrates the benefit of the q-group index being tailored toward the GPU architecture.

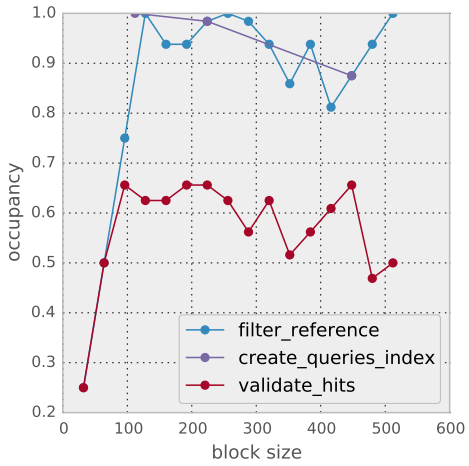

Figure 3: The occupancy of GPU cores depending on the thread block size. Shown are representative patterns for OpenCL kernels from the three main steps of the algorithm: building the index (create_queries_index), filtration (filter_reference) and validation (validate_hits).

Sensitivity. To assess the sensitivity of the algorithm we use the Rabema [5] benchmark that allows to compare mapping results based on a formalized framework. First, 1000 Illumina[3] reads of length 100 were simulated using Mason [4] with the Saccharomyces cerevisiae genome[4], default parameters and error rates. Second, the simulated reads were mapped to the genome using RazerS 3 [23] with full sensitivity and different error tolerances. The mapped reads were used to generate gold standards for Rabema to test against.

We analyze the sensitivity of our algorithm using q-grams of length 16, because it is computationally optimal on the current GPU hardware. We configure PEANUT to provide all semi-global alignments of a read and leave all other parameters at their default values. Sensitivity is assessed by the Rabema measure "Normalized found intervals" [5] and all alignments of a read are considered (all-mode). With a percent identity threshold of 60 (see Section 3.3) our algorithm provides 100% sensitivity for error rates below 7%, at least 99.77% sensitivity for error rates up to 10% and still 98.96% sensitivity with an error rate up to an unrealistically high 20%. With a stricter threshold of 80, PEANUT still reaches 98.89% sensitivity for the latter.

In general, we advise to set the percent identity threshold to be slightly more permissive than the expected error rate. This is because the replacement of N-characters in the reads and the reference with random bases can introduce additional mismatches. Above rates are far better than the worst case sensitivity that can be expected by applying the pigeonhole principle (i.e., with reads of length 100 and q-grams of length 16, we can expect to find at least one perfectly matching q-gram for all alignments with at most 6 errors), such that using 16-grams appears to be a reasonable default choice in practice.

Running time vs. sensitivity. We evaluate the run time performance using a dataset of 5 million simulated Illumina HiSeq reads (created with Mason [4] at default error rates) on the ENSEMBL human reference genome[5] version 37. Run time performance was measured three times as the overall wall clock time and the best result was selected (best of 3). The sensitivity of the alignments (i.e. the percentage of rediscovered alignments from the simulation) was assessed by using the Rabema

[3] `http://www.illumina.com`
[4] as provided by the Rabema data package, `http://www.seqan.de/projects/rabema/`, October 2013
[5] `ftp://ftp.ensembl.org/pub/release-74/fasta/homo_sapiens/dna/`

Table 1: Performance and sensitivity of PEANUT and other state of the art read mappers on the human reference genome with (a) 5 million simulated Illumina reads of length 100, (b) 5 million real Illumina HiSeq 2000 reads from the human exome and (c) 10 million paired-end real Illumina HiSeq 2000 reads from the human exome. For the simulated dataset, sensitivity was measured with Rabema [5]. Running time is given for the quickest run of three repetitions.

dataset	mapper	type	time [min:sec]	sens. [%]
a	PEANUT	best-stratum	**1:55**	**98.62**
a	BWA-MEM	best	3:16	96.99
a	Bowtie 2	best	5:21	96.85
a	PEANUT	all	**18:29**	98.74
a	RazerS 3	all	199:55	**98.83**
b	PEANUT	best-stratum	**1:33**	N/A
b	BWA-MEM	best	1:58	N/A
b	Bowtie 2	best	3:12	N/A
b	PEANUT	all	**10:52**	N/A
b	RazerS 3	all	89:38	N/A
c	PEANUT	best-stratum	**3:08**	N/A
c	BWA-MEM	best	4:44	N/A
c	Bowtie 2	best	8:18	N/A
c	PEANUT	all	**21:54**	N/A
c	RazerS 3	all	150:59	N/A

measure "Normalized found intervals" calculated on the simulated data. Rabema has been run in "oracle mode" with the "any-best" benchmark type [5], using the original position information from the read simulation. This means, that the benchmark counts each alignment that is equally good as the one from the read simulation as correct. We consider this benchmark to be fair, since the best mappers BWA-MEM and Bowtie 2 per default only provide a single alignment even if there are multiple ones of the same quality.

The benchmark was conducted on an Intel Core i7-3770™ system (4 cores with hyperthreading, 3.4GHz, 16 GB RAM), an NVIDIA Geforce™ 780 GPU (12 SMs, 3GB RAM) and a 7200 rpm hard disk. We evaluated two modes of PEANUT. First, PEANUT was configured to find the best stratum of semi-global alignments (best-stratum mode) for each read. Second, PEANUT was configured to find all semi-global alignments (all mode) for each read. Further, we benchmarked the newest generation of BWA (BWA-MEM, [10]), Bowtie 2 [9] and RazerS 3 [23]. (At the time of writing, we were not able to obtain a working binary or source distribution of the GPU based read mappers SOAP3 [13] and SOAP3-dp [14] for our setup. Further, NextGenMap [21] was not able to detect our OpenCL setup, and therefore refused to run.) While BWA-MEM and Bowtie 2 provide only the best hit, RazerS 3 per default reports all possible alignments. All read mappers were configured to output alignments in SAM format [12] directly to the hard disk.

Table 1 compares PEANUT with its competitors. First, PEANUT in best-stratum mode outperforms the best-mappers BWA-MEM and Bowtie 2 in terms of speed and sensitivity. Second, PEANUT in all-mode is 10 times faster than the all-mapper RazerS 3, while maintaining a comparable sensitivity. While Bowtie 2 provides an all-mode, too, it did not terminate in competetive time due to extensive memory requirements exceeding the capabilities of our test system.

In addition, we measured the run times using 10 million paired-end reads of of the publicly

Table 2: Performance of PEANUT on secondary test system with a four years old Geforce 580 GPU vs. a more recent Geforce 780. Times are in min:sec. See also Table 1.

dataset	type	time (580)	time (780)	difference
a	best-stratum	2:11	1:55	13.9%
a	all	20:56	18:29	13.3%
b	best-stratum	1:45	1:33	12.9%
b	all	11:49	10:52	8.7%
c	best-stratum	3:28	3:08	10.6%
c	all	23:55	21:54	9.2 %

available exome sequencing dataset ERR281333[6] [17] obtained from an uveal melanoma sample sequenced with Illumina HiSeq 2000. When interpreting the 5 million forward direction reads of above dataset as single-end reads, the run time results resemble those of the simulation benchmark, but all four mappers perform slightly faster. This is likely caused by the greater locality of data structure accesses caused by the exome-sequenced reads, which should be less distributed across the genome. Nevertheless, PEANUT is again faster than its competitors in both modes. Finally, we investigated the paired-end performance of all mappers using above dataset. Again, PEANUT proves to be faster than the competitors.

While PEANUT requires at least 2.5 GB of GPU memory for filtration and validation, it is not restricted to running on high-end GPU models like the Geforce 780 used above: Table 2 shows that a speedup can be maintained when benchmarking on a different test system with an Intel Core i7-2600 (3.4GHz, 16GB RAM) and a four years old NVIDIA™ Geforce 580 GPU. We see that the older system is 9% to 14% slower, depending on the dataset and still faster than the best competitor in Table 1.

Hit-rank vs. sensitivity. Finally we evaluate the ability of the hit-rank (see Section 3.4) to distinguish between true and false positives. Here, true positives are the correctly identified true sampling positions of the reads, whereas false positives are reported mapping locations that may have the same alignment score but are not the true origins of a read. For this purpose the sensitivity was calculated for each hit-rank in the first 100000 reads of the simulated human dataset from above. Figure 4 shows the accumulated sensitivity vs. the considered hit-ranks. As can be seeen, the hit-rank (and hence the mapping qualities of PEANUT) provides a sharp distinction between true and false positives and small ranks account for most of the sensitivity.

5 Discussion

We presented PEANUT, a read mapper exploiting the capabilities of modern GPUs using OpenCL. PEANUT is based on the q-group index, a novel variant of a q-gram index that has a particularly small memory footprint and is therefore tailored for modern GPUs. To the best of our knowledge, the presented datastructure and algorithm is the first implementation of the filtration-validation approach that works completely on the GPU. The benefit of this is illustrated by PEANUT showing supreme speed over state of the art read mappers like BWA, Bowtie 2 and RazerS 3. This holds for both the high-end GPU NVIDIA™ Geforce 780 and a four years old NVIDIA™ Geforce 580. Thereby, PEANUT maintains a high sensitivity that is comparable to RazerS 3 and higher than BWA and Bowtie 2.

[6] `http://www.ebi.ac.uk/ena/data/view/PRJEB3979`

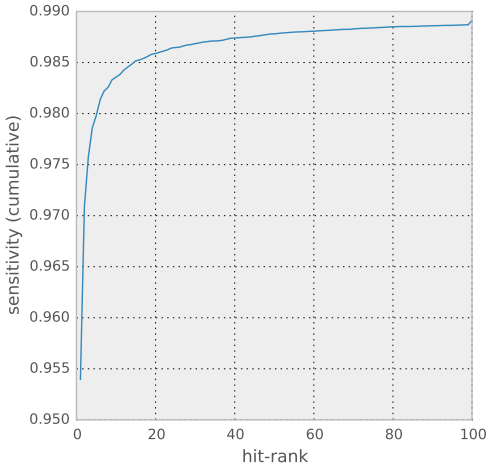

Figure 4: Hit-rank vs. sensitivity for 100000 simulated reads from the human reference genome. For each hit-rank k, the y-axis depicts the sum of the sensitivities up to k.

PEANUT is distributed under the MIT license as an open source Python[7] software package. Filtration and validation where implemented in OpenCL, using the PyOpenCL package [7] and postprocessing was implemented in Cython [1]. Documentation and installation instructions are available at `http://peanut.readthedocs.org`, where we further provide a Snakemake [8] workflow of all analyses conducted in this work.

Acknowledgments. Part of this work was funded by the German Research Foundation (DFG), Collaborative Research Center (Sonderforschungsbereich, SFB) 876 "Providing Information by Resource-Constrained Data Analysis" within project C1, see `http://sfb876.tu-dortmund.de`. SR acknowledges support from the Mercator Foundation within project MERCUR Pe-2013-0012 (UAMR Professorship "Computational Biology").

References

[1] Stefan Behnel, Robert Bradshaw, Craig Citro, Lisandro Dalcin, Dag Sverre Seljebotn, and Kurt Smith. Cython: The best of both worlds. *Computing in Science and Engg.*, 13(2):3139, March 2011.

[2] Guy E. Blelloch. *Vector Models for Data-parallel Computing*. MIT Press, Cambridge, MA, USA, 1990.

[3] Jochen Blom, Tobias Jakobi, Daniel Doppmeier, Sebastian Jaenicke, Jörn Kalinowski, Jens Stoye, and Alexander Goesmann. Exact and complete short-read alignment to microbial genomes using graphics processing unit programming. *Bioinformatics*, 27(10):1351–1358, May 2011.

[4] Manuel Holtgrewe. Mason – a read simulator for second generation sequencing data, October 2010. Technical Report TR-B-10-06, Institut für Mathematik und Informatik, Freie Universität Berlin.

[7] `http://www.python.org`

[5] Manuel Holtgrewe, Anne-Katrin Emde, David Weese, and Knut Reinert. A novel and well-defined benchmarking method for second generation read mapping. *BMC Bioinformatics*, 12(1):210, May 2011.

[6] Heikki Hyyrö. A bit-vector algorithm for computing levenshtein and damerau edit distances. *Nordic Journal of Computing*, page 2003, 2003.

[7] Andreas Klöckner, Nicolas Pinto, Yunsup Lee, Bryan Catanzaro, Paul Ivanov, and Ahmed Fasih. PyCUDA and PyOpenCL: a scripting-based approach to GPU run-time code generation. *Parallel Computing*, 38(3):157–174, March 2012.

[8] Johannes Köster and Sven Rahmann. Snakemake – a scalable bioinformatics workflow engine. *Bioinformatics*, 28(19):2520–2522, October 2012.

[9] Ben Langmead and Steven L. Salzberg. Fast gapped-read alignment with bowtie 2. *Nature Methods*, 9(4):357–359, April 2012.

[10] Heng Li. Aligning sequence reads, clone sequences and assembly contigs with BWA-MEM. arXiv e-print 1303.3997, March 2013.

[11] Heng Li and Richard Durbin. Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics*, 25(14):1754–1760, July 2009.

[12] Heng Li, Bob Handsaker, Alec Wysoker, Tim Fennell, Jue Ruan, Nils Homer, Gabor Marth, Goncalo Abecasis, and Richard Durbin. The sequence Alignment/Map format and SAMtools. *Bioinformatics*, 25(16):2078–2079, August 2009.

[13] Chi-Man Liu, Thomas Wong, Edward Wu, Ruibang Luo, Siu-Ming Yiu, Yingrui Li, Bingqiang Wang, Chang Yu, Xiaowen Chu, Kaiyong Zhao, Ruiqiang Li, and Tak-Wah Lam. SOAP3: ultra-fast GPU-based parallel alignment tool for short reads. *Bioinformatics*, 28(6):878–879, March 2012.

[14] Ruibang Luo, Thomas Wong, Jianqiao Zhu, Chi-Man Liu, Xiaoqian Zhu, Edward Wu, Lap-Kei Lee, Haoxiang Lin, Wenjuan Zhu, David W. Cheung, Hing-Fung Ting, Siu-Ming Yiu, Shaoliang Peng, Chang Yu, Yingrui Li, Ruiqiang Li, and Tak-Wah Lam. SOAP3-dp: fast, accurate and sensitive GPU-Based short read aligner. *PLOS ONE*, 8(5):e65632, May 2013.

[15] Santiago Marco-Sola, Michael Sammeth, Roderic Guig, and Paolo Ribeca. The GEM mapper: fast, accurate and versatile alignment by filtration. *Nature Methods*, 9(12):1185–1188, December 2012.

[16] Tobias Marschall, Ivan G. Costa, Stefan Canzar, Markus Bauer, Gunnar W. Klau, Alexander Schliep, and Alexander Schönhuth. CLEVER: clique-enumerating variant finder. *Bioinformatics*, 28(22):2875–2882, November 2012.

[17] Marcel Martin, Lars Mahöfer, Petra Temming, Sven Rahmann, Claudia Metz, Norbert Bornfeld, Johannes van de Nes, Ludger Klein-Hitpass, Alan G. Hinnebusch, Bernhard Horsthemke, Dietmar R. Lohmann, and Michael Zeschnigk. Exome sequencing identifies recurrent somatic mutations in EIF1AX and SF3B1 in uveal melanoma with disomy 3. *Nature Genetics*, 45(8):933–936, August 2013.

[18] Gene Myers. A fast bit-vector algorithm for approximate string matching based on dynamic programming. *J. ACM*, 46(3):395415, May 1999.

[19] Kim R Rasmussen, Jens Stoye, and Eugene W Myers. Efficient q-gram filters for finding all epsilon-matches over a given length. *Journal of computational biology: a journal of computational molecular cell biology*, 13(2):296–308, March 2006.

[20] Adam Roberts and Lior Pachter. Streaming fragment assignment for real-time analysis of sequencing experiments. *Nature Methods*, 10(1):71–73, January 2013.

[21] Fritz J Sedlazeck, Philipp Rescheneder, and Arndt von Haeseler. NextGenMap: fast and accurate read mapping in highly polymorphic genomes. *Bioinformatics (Oxford, England)*, 29(21):2790–2791, November 2013.

[22] Cole Trapnell, Lior Pachter, and Steven L. Salzberg. TopHat: discovering splice junctions with RNA-Seq. *Bioinformatics*, 25(9):1105–1111, May 2009.

[23] David Weese, Manuel Holtgrewe, and Knut Reinert. RazerS 3: Faster, fully sensitive read mapping. *Bioinformatics*, 28(20):2592–2599, October 2012.